

GLOSSARY

\$@ — A special shell variable that includes all of the parameters on the command line.

.bashrc — A configuration script that is executed each time the user starts a `bash` shell.

.profile — A configuration script that can be located in each user's home directory.

/etc/bashrc — A script that is executed each time any user on the system starts a `bash` shell. This script is not included by default on all Linux distributions, but can be created if needed.

/etc/fstab — Configuration file that contains a file system table with devices, mount points, file system types, and options. Used by the `mount` command.

/etc/group — Configuration file in which group information (group names and membership lists) is stored.

/etc/passwd — Configuration file in which user account information is stored.

/etc/profile — A script containing configuration information that applies to every user on the Linux system.

/etc/shadow — Configuration file in which encrypted user passwords and password configuration data are stored.

/etc/skel — Directory containing files that will be used to populate a new user's home directory at the time it is created.

/proc file system — A method of viewing what the operating system kernel is doing at all times, using a section of the Linux directory structure and a file system interface (using the same commands used for regular files).

AC power — The standard alternating current power coming into a building from a public utility; the power from a wall socket.

action — A field in the `syslog.conf` configuration file that determines what to do with messages matching the selector on that line.

active partition — The partition that receives control from the BIOS when the system is turned on.

alias — A string of characters that the shell substitutes for another string of characters when a command is entered.

alias — Command used to create a text substitution in a command-line shell, effectively giving any Linux command a new name.

application — A program (such as a word processor or spreadsheet) that provides a service to a person using the computer, rather than simply managing the computer's resources.

apropos — Linux command used to show all man pages that contain a keyword.

array — A collection of multiple hard disks. *See RAID.*

at — A command that lets you enter one or more commands to be executed once at some future time.

atq — A command that lists each of the jobs that have been submitted using the `at` command, with a job number and the date and time when the job will be executed.

atrm — A command that deletes (removes) a job from the queue used by `atd` to execute commands.

authentication — The process of identifying a user to a computer system via some type of login procedure.

awk — A programming language that developers use to create scripts for working on text files and completing other complex tasks.

backup — A copy of data on a computer system.

back-up level — A definition of how much data is to be backed up in comparison with another back-up level. When performing a back-up operation at a given back-up level, all of the data that has changed since the last backup of the previous level is recorded.

back-up media — A device where data can be stored, such as a tape cartridge, writeable CD, or even a floppy disk.

back-up plan — A written document that outlines when, how, and, perhaps, why various files and file systems will be backed up, stored, and—when necessary—restored to prevent permanent data loss.

background application — An application that does not stop the program that started it from going on to other tasks.

bang — In Linux jargon, an exclamation point character.

bash — Short for *Bourne Again shell*, an enhanced and extended version of the Bourne shell created by the GNU project for use on many UNIX-like operating systems. **bash** is the default Linux shell.

batch — A command that executes scheduled tasks when the system load average drops below 0.8.

benchmark program — A program that provides a numeric measurement of performance for part of a system.

Beowulf — A cluster of multiple Linux servers operating in parallel as a supercomputer to solve complex problems.

bg — Command used to place a job (process) in the background (either by suspending it or by preventing its output from appearing in the current shell's terminal window), thus allowing the shell prompt to become active again.

binary code — Machine-readable instructions used to execute a program.

binary file — *See* executable file.

BIOS (Basic Input/Output System) —

Information stored in ROM that provides instructions to the operating system for using the devices on a computer.

block — A unit of storage on a file system. A standard block contains 1024 characters (bytes), or two sectors.

block special file — A type of file (normally located in `/dev`) referring to a physical device that transfers data in blocks of characters, such as a hard disk drive.

bonnie — A hard disk benchmarking program.

boot disk — A disk used to launch the Linux operating system stored on your hard disk. It can be used in normal operating situations to start the Linux system. A special version of the boot disk launches the Linux installation program.

boot manager — A program that lets you select an operating system to use each time you boot the computer.

boot parameter — A piece of information passed directly to the Linux kernel as the system is being booted. These parameters are normally used to affect how Linux recognizes hardware devices or to enable certain features of the operating system.

boot record — A small area on each partition that contains a program to launch the operating system on that partition.

bootable CD-ROM drive — A CD-ROM drive that can launch an operating system (or other program) directly from a CD without accessing the hard disk. (This feature of the CD-ROM drive must be enabled by the BIOS.)

bottleneck — Part of a computer system that slows down completion of the task at hand.

bounce queue (bq) — An option code within the `printcap` file that causes print jobs to be processed by the local print filter before being sent to a remote printer.

Bourne shell — The original shell for UNIX, written by Stephen Bourne.

broadcast address — A special IP address that sends a packet of data to all computers on the local network.

buffers — Areas of memory dedicated to holding the working data associated with a running program.

byte — Space within a computer system sufficient to store one character.

C shell — A shell developed by Bill Joy in the 1970s. He focused on adding easy-to-use features for interactive work at the shell prompt. (Most of these features were later added to the `bash` shell as well.) The C shell is not popular for shell programming because its syntax is more complex than that of the Bourne, `bash`, and Korn shells.

cached data — Information from the hard disk that is stored temporarily in RAM, under the assumption that the data will soon be needed by an application. Caching data improves system performance, because data can be accessed faster from RAM than from the hard disk.

caching — The process of storing data from the hard disk in RAM so that it can be accessed more rapidly (because RAM is much faster than a hard disk).

cat — Command used to dump the contents of a file to STDOUT.

cd — Command used to change the directory you are working in (the current working directory).

character special file — A type of file referring to a physical device (such as a serial port) that transfers data in single characters.

chief information officer (CIO) — The executive in an organization who determines how information systems are used within the organization to further its goals or mission effectively.

chmod — Command used to change the file permissions assigned to a file or directory.

chown — Command used to change the ownership of a file or directory.

COAS (Caldera Open Administration System) — A set of graphical utilities developed by Caldera Systems and used to manage many aspects of a Linux system.

command interpreter — (More commonly called a shell in Linux.) A command-line environment in which a user can enter commands to be launched.

command-line option — A command-line parameter.

command-line parameter — An additional piece of information (besides the program name) that is included on the command line when a program is started.

command-line window — A window within a graphical environment that permits you to enter commands at the keyboard.

comment — A line in a script that begins with a `#` character. Comments are not processed by the shell, but are only included to help someone reading the file understand the purpose of the script or how it functions.

Common Gateway Interface (CGI) — A method of communication between two programs using the standard input and standard output channels.

compiled language — A computer language for which the source code is converted to a binary file long before the program will be run by users.

compiler — A special program that converts the source code of a compiled language into a binary file.

computer language — A set of words and syntax rules that can be arranged in predefined ways to cause a computer to perform tasks defined by the person using the language.

cooperative multitasking — A technique in which an operating system kernel must wait for a program to yield control to other programs.

copyleft — An ironic term that refers to the GNU General Public License (the GPL), signifying a radical departure from standard copyright.

cp — Command used to copy a file or directory from one location or name to another.

cpio — A Linux archiving program. The `cpio` command also reads archive files created by the `tar` command.

cron job — A command or script that you have scheduled to be executed at a specific time in the future.

crontab — A command that lets you enter one or more commands to be executed repeatedly at intervals that you designate.

current working directory — The directory in which you are working.

cylinder — A set of tracks at the same location on all the platters of a hard disk.

daemon — A background process that does not display status information on the screen. Instead, daemons normally write information to log files.

default shell — The default command-line interpreter used in most Linux systems (`bash`).

defragmenting — The process of rearranging the files on a file system so that all the parts of a file are located next to each other on the physical hard disk. Defragmenting a hard disk increases system performance and reduces wear on the storage device.

desktop environment — A graphical application that provides a comprehensive interface, including system menus, desktop icons, and the ability to easily manage files and launch applications.

device drivers — Software that provides access to additional hardware, beyond core device support provided by the kernel.

df — Short for display file systems. Command used to display file system summary information such as device, mount point, percentage used, and total capacity.

direct memory access (DMA) channel — A communication method within a computer that allows a device to read and write directly to the computer's RAM, without going through the microprocessor first.

directory record — A file containing the names and inode numbers of other files.

disaster plan — An organized written plan that describes how to respond to various threats to an information system such as Linux.

disk image — A single file that contains an exact copy of a floppy disk.

disk optimization — *See* defragmenting.

dmesg — Program that displays the contents of the kernel ring buffer. This buffer normally contains hardware configuration data generated during system start-up.

DNS server — A computer that uses the DNS protocol to convert domain names and hostnames to IP addresses.

do — A command used with the `done` command to enclose a block of commands to be executed by a `for` command.

domain name — A name assigned to a network.

Domain Name Service (DNS) — A protocol that maps human-readable domain names and hostnames to IP addresses that correspond to networks and individual computers.

done — A command used with the `do` command to enclose a block of commands to be executed by a `for` command.

DOS — An operating system developed for personal computers in about 1980. It gained widespread acceptance when IBM introduced the first IBM PC.

downtime — Occasions when an organization's computer systems cannot respond to requests for information.

driver — A software program that provides abstract services for a hardware component, such as opening files or reading character input.

du — Short for disk usage. Command used to display disk space used by a directory and each of its subdirectories.

dual-boot system — A computer that allows a person to choose which operating system to start each time the computer is booted (turned on).

dumpe2fs — A utility used to display technical statistics and parameters about a Linux `ext2` file system.

duplexing — Term used to describe a system in which the contents of two file systems, that are located on different hard disk controllers, contain identical information. Compare to "mirroring," a technique that provides identical information on two file systems but without redundant disk controllers.

Dynamic Host Configuration Protocol (DHCP)

— A protocol that allows a computer to obtain an IP address dynamically from a network server at the time the computer is turned on.

dynamically linked applications — Linux programs that do not include the library functions that they require in order to operate. The libraries must be installed (as shared libraries) on the Linux system on which the applications are executed.

e2defrag — The utility used to defragment a standard Linux `ext2` partition.

echo — A command that prints text to the STD-OUT channel—to the screen unless output has been redirected.

else — A command that extends the capability of an `if/then` statement by adding a block of commands that are only executed if a test returns a value of false (that is, if the test fails).

end user — An individual who uses the computer systems in an organization to accomplish assigned tasks, but relies on a system administrator to keep those systems running smoothly.

environment variables — Set of named values (name-value pairs) that provide information to programs running in a user's environment.

executable file — A file containing numeric codes that a computer can execute. Created from a source code file by a compiler, the executable file is the program that a user can run.

execute permission — A file permission that allows a user to launch a file as a program or see a file within a directory. Represented by a letter `x`.

exit — A command that stops the execution of a script immediately.

export — Command used to make a newly created environment variable available to other programs running in the same environment.

ext2 — The default file system type for Linux.

facility — A category assigned to a system message, identifying the type of program providing the message.

FAT32 — The default file system type for Windows 98.

fault tolerance — The condition of being able to tolerate errors or events that might otherwise cause system failure.

fdformat — Command used to format a floppy disk.

fdisk — A utility used to create hard disk partitions and configure how they are used.

fg — Command used to bring a job (process) running in a shell to the foreground so that the job controls the shell's terminal window.

fi — A command that marks the end of an `if/then` statement.

file — Command used to print a summary of the type of data contained in a file.

file extension — The last part of a filename after a period.

file handle — An internal storage mechanism that allows a single file to be opened and used in Linux.

file manager window — A graphical window that displays the contents of a directory (usually as a collection of icons) and lets you work with the files and directories using menus, mouse clicks, and dialog boxes.

file permissions — Codes that define the type of access that a user has to a file or directory on the Linux system.

file record — An information item within an `ext2` file system that includes a filename and an inode number. The inode itself contains detailed information about the file.

file system — A collection of data, normally stored on a device such as a hard disk partition, which can be accessed in Linux via the directory structure.

filtering — The process of adding, removing, or altering data in the text file based on complex rules or patterns.

find — Linux command used to search the file system for files matching certain characteristics.

FIPS — A program that creates two separate partitions from an existing Windows partition.

for — A command that repeats a block of commands one time for each item in a list that you provide.

for loop — A block of commands that is repeatedly executed according to the parameters provided with the `for` command.

free — Linux command used to display the amount of free and used memory (physical and virtual), with basic information about how that memory is being used.

Free Software Foundation (FSF) — An organization founded by Richard Stallman to promote his ideals of freely available software and to create and distribute that software.

fsck — A utility used to check the integrity of an `ext2` file system.

fully qualified domain name (FQDN) — The hostname of a computer and the domain name of the network to which the computer is attached.

function — A set of computer programming code that completes a certain task for a program.

function library — A file containing a collection of commonly used functions that any program can use as it runs.

gateway address — The IP address of the computer on a local network that can send packets of data outside that network.

gcc — A C language compiler. Probably the best known product of the GNU project.

gigabyte (GB) — A measure of space on computers equal to 1024 megabytes, or roughly enough space to store 1 billion characters.

Gnome System Monitor — A graphical utility for the Gnome Desktop that is used to monitor and control processes running on Linux. Also called `gtop`. Similar to the `kpm` program for KDE.

GNU General Public License (GPL) — The free software license that Richard Stallman of the Free Software Foundation developed for the programs created by the GNU project.

GNU project — An effort by the Free Software Foundation to create a free UNIX-like operating system. Much of a Linux distribution comes from the GNU project.

graceful shutdown — The technique of stopping all Linux services and shutting down all file access in an orderly way before turning off or rebooting the computer.

graphical interface — Software that provides mouse-driven applications with menu bars, buttons, and so forth.

graphical libraries — Collections of programming functions that an X client can use to more efficiently create and manage the elements of a graphical environment.

grep — Linux command used to search within files for lines containing a certain pattern.

groff — A command used to format and display documents that are created using roff mark-up codes.

group — A named account that consists of a collection of users. Each member of a group has access to files owned by that group.

groupadd — Command used to add a new group to a Linux system.

group permissions — A set of three file permissions (r, w, and x) that apply to members of the group assigned to a file or directory.

gunzip — Command used to uncompress a file that has been compressed using gzip.

gzip — Command used to compress any file on a Linux system.

halt — Command used to shut down all services and then stop the computer with the message “System halted.”

hard disk — Magnetic storage space for data such as the operating system and data files.

hard link — A pointer to an inode that is already pointed to by at least one other file record.

hard wired — Computer functionality that is arranged in the wires and other components that make up a computer. Hard-wired functionality cannot be easily altered.

hardware-based RAID — A RAID array that is contained in a separate hardware device (a RAID subsystem) and is controlled by a CPU and other components separate from the CPU of the Linux system.

Help Desk — A department in many organizations that assists end users in solving problems related to information technology.

hexadecimal (hex) — A numbering system using base-16. Hex uses 0 to 9, plus the letters A through F (usually capitalized) to count the numbers 10 through 15.

high availability — Term used to refer to the processes, products, or programs involved in ensuring that a system experiences as little downtime as possible. The goal for all high availability systems is 100% uptime.

high availability cluster — A group of servers that process the same tasks (resource groups) and take over each others' functionality in the event of an outage or failure.

history — A command used to display all of the stored commands in the history list.

history feature — A feature of the shell that records in a list (the history list) each of the commands that you enter at the shell prompt.

history list — A list that contains the most recently executed commands. (Normally at least 100 commands are included in the history list.)

home directory — The location where all of a user's personal files are stored.

host — A computer attached to a network.

hostname — A single word used to name a computer.

hot-swapping — The process of removing and replacing a failed hard drive from a RAID hardware device or specialized server without turning off the power to the device.

HOWTOs — Documents within the Linux Documentation Project that cover specific topics.

IDE — A low-cost, easy-to-manage interface used by most new computers to connect hard disks and CD-ROM drives to the CPU.

IDE controller — A computer hardware component used to communicate between an IDE-compatible hard disk or other IDE device and the microprocessor.

if — A command used to introduce a test within a shell script. An **if** command is always followed by a **then** command.

if/then/else statement — A set of commands used to determine whether other commands in a script are executed. An **if/then/else** statement is one kind of selection statement.

index — A counting variable used within a loop statement. The index acts as a marker to count how many times the loop has executed a list of commands.

info — Linux command used to access online command reference information.

Information Systems Department (IS) — The area of an organization in which the staff are responsible for maintaining computer and information systems that support the employees in their work (also called the IT Department in some organizations).

Information Technology Department (IT) — See Information Systems Department (IS).

init — A daemon that acts as a control process to start the first processes on Linux, such as the login screens. Also, a command used to switch the system to a different run level.

inode — A file information record, identified by a unique number within a file system, which contains detailed information about a block of data commonly called a file.

insmod — Command used to copy a module file from the hard disk and add it to the Linux kernel running in memory.

install disk — A disk used to start the Linux installation program on some distributions of Linux. See boot disk.

installation source — The set of files from which Linux is installed. These files are normally stored on a Linux CD.

installation type — A specification indicating which Linux software to install; the correct installation type depends on how the Linux system will be used.

interrupt request (IRQ) — A numbered channel of communication allowing a device to inform the system that some action needs to be taken for the device.

interpreted language — A language for which the source code is converted to numeric codes at the time a user runs the program. This conversion takes place each time the user runs the program. For this reason, interpreted languages are generally much slower than compiled languages.

interpreted program — A computer program that is converted from human-readable form to a format that can be used by a computer (numeric codes) at the moment you execute the program.

interpreter — A special program that converts the source code of an interpreted language into numeric codes that a computer can execute.

I/O ports — Memory addresses used by a device for memory-mapped I/O.

IP — A networking protocol used to send packets of information across a network connection.

IP address — An identifying number assigned to a computer or device that uses IP to communicate across a network.

iteration — An occurrence of an event or process that can or must be done many times.

jobs — Command used to list jobs (processes) started in the current shell environment.

jukebox — A back-up device that holds multiple back-up media (such as multiple tape cartridges or writeable CDs) and that can switch between them without assistance from a system administrator.

kernel — The core of the operating system, which interacts directly with the computer hardware.

kernel modules — Files containing computer code that can be loaded into the kernel or removed from the kernel as needed.

kernel ring buffer — A small area of memory that holds internal kernel messages. These messages can be viewed using the `cat /proc/kmsg` utility.

keywords — Words used in a computer language to define a specific task or meaning.

kill — Command used to send signals to processes, often to end them via a SIGTERM or SIGKILL signal.

klogd — A background program (or daemon) used to log kernel messages according to the configuration given in the `syslog.conf` configuration file.

Korn shell — A revision of the Bourne shell that includes the interactive features of the C shell but that maintains the Bourne shell programming style. The Korn shell was written by David Korn. It is available in Linux as the Public Domain Korn shell, `pdksh`.

kpm — The KDE process manager, a graphical utility for the KDE graphical environment that is used to control Linux processes. Similar to the Gnome System Monitor.

ktop — A KDE-based graphical utility used to manage Linux processes. Similar to the `kpm` program.

LaTeX — A version of the mark-up language TeX that includes numerous macros for easy document creation.

ldd — Command used to list the function libraries that a program uses.

legacy systems — Computer systems that an organization already owns. This term usually refers to systems that are no longer state of the art.

less — Command used to print the contents of a file one screenful at a time. It allows you to move around

in the file and otherwise control the command by using the keyboard.

LGPL — A special version of the GNU General Public License intended to govern both free and commercial software use of software libraries.

lilo — Command used to read the `lilo.conf` configuration file and update the hard disk boot information based on the configuration.

LILO (Linux Loader) — The boot manager included with Linux.

link — A special file record that refers to the same physical file data as another file record.

LinuxConf — Graphical configuration and administration utility for Linux, developed and supported by Red Hat Software.

Linux distribution — A Linux operating system product that includes the Linux kernel plus many software components, installation tools, documentation, and so forth.

Linux Documentation Project (LDP) — One of the first efforts to document how Linux is used. Started by Matt Welsh.

lizardx — A graphical utility used for configuring the X Window System.

ln — Command used to create a symbolic or hard link.

locate — Linux command used to search an index of the file system for items matching a given pattern.

log file — File that contains detailed records of activity on a Linux system.

logger — A program that lets you send a message to the `syslog` function. Such messages are written to the log files according to the configuration in `syslog.conf`.

logging in — The process of identifying yourself as a valid user who has been assigned a certain set of access rights.

logrotate — A program that manages the rotation of multiple log files at regular intervals according to a configuration file created by the system administrator.

logscanner — A log analysis program available for download at http://xmmission.linuxberg.com/conhtml/adm_log.html.

LogWatch — A log analysis program available for download at http://xmmission.linuxberg.com/conhtml/adm_log.html.

loop statement — A statement used to determine whether part of a computer program should be executed more than once.

lpd — Short for *line printer control*. The Linux printer control utility.

lpd — The line printer daemon, which sends files prepared by **lpr** to the physical printer device or remote print server.

lpq — A utility that lists each print job in a print queue and includes information such as the owner and size of each job.

lpr — A command that prepares files to be sent to a physical printer device, effectively “printing” files for Linux users.

lprm — A command that deletes a print job from a print queue.

ls — Command used to list the files in a directory.

lsmod — Command used to list the modules that are installed in the Linux kernel.

Macintosh — A computer developed by Apple Computer that integrated the operating system and the graphical interface.

macro — A set of commands that can be executed as one by referring to the name of the macro.

magic filter — A print filter program that automatically processes a file into the correct output format based on the file’s type.

make — A command that uses information stored in a configuration file to compile source code into a binary program, or otherwise prepare a program to be executed.

man pages — Online reference pages for Linux commands. The man pages are accessed using the **man** command.

mark-up languages — Computer languages that define a series of codes indicating how to format a document.

Master Boot Record (MBR) — A small area of the first hard disk that contains a program to decide how to start an operating system.

megabyte (MB) — A measure of space on computers equal to 1,048,576 bytes, or enough space to store roughly 1 million characters.

memory-mapped input/output (memory-mapped I/O) — A technique that assigns a range of memory addresses in a computer as a place for a device to send and receive data.

message — A description of what is happening within a program.

messages — The main system log file in Linux, usually stored in the directory **/var/log**.

Microsoft Windows — The leading graphical interface for DOS.

mirroring — Term used to describe a system in which the contents of two file systems contain identical information. Mirroring improves data access speed and provides fault tolerance in the event that one of the file systems fails.

mkdir — Command used to create a new directory.

mke2fs — Command used to format a device such as a hard disk partition with an **ext2** file system.

mkfs — Command used to format devices using various file system types. The **ext2** default type for Linux file systems can be indicated as an option. *See also mke2fs.*

mkswap — Command used to format a partition as a swap space for the Linux kernel.

modal editor — A text editor that uses multiple modes for editing text and entering commands to apply to that text.

modprobe — Command used to load a module with all of its required supporting modules.

module parameters — Information needed by a module to locate system resources. The parameters are added after the module name when using the `insmod` or `modprobe` command.

more — Command used to print the contents of a file one screenful at a time. The `more` command is similar to the `less` command but with fewer keyboard control options.

mount — Command used to make a logical or physical device available as a file system in the Linux directory structure.

mount point — The place or path in the Linux directory structure where a file system is accessed.

multithreading — A technique used within multiprocessor operating systems to divide a larger task between multiple processors.

multiuser system — An operating system on which numerous users can log in to the same computer (usually over a network connection).

mv — Command used to rename a file or directory or move it to a new location.

nbench — A benchmark program developed by *BYTE* magazine and ported to Linux.

nesting — A programming method in which one selection or loop statement contains another selection or loop statement.

network mask — A set of numbers that tells the networking system in Linux how to identify IP addresses that are part of the local network.

nice — Linux command used to set the nice level of a Linux process as it is being launched.

nice level — The priority level assigned to a Linux task.

NTFS — The default file system type for Windows NT.

OpenSource — A trademarked name often used to refer to software licensed under the GPL.

operating system — Software that provides a set of core functionality for other programs to use in working with the computer hardware and interfacing with the user running the computer.

other permissions — A set of three file permissions (r, w, and x) that apply to all users on the Linux system who are not the owner of the file or directory in question and are not members of the group assigned to the file or directory.

pac — A command that displays reports based on print accounting data. It stands for *print accounting*.

page — A block of 4 KB of memory. A page is the unit of memory in which the Linux kernel moves data to and from the Linux swap partition.

page description language — A special set of codes that determine the graphics elements, text font, and everything else about how information appears on a printed page.

parent directory — The directory that is one level above the current directory.

parity — A technique that allows corrupted data to be reconstructed using an extra piece of information (the parity information) that is created as the data is stored. Parity information provides redundancy for a piece of information.

parity stripe — Parity information stored as part of a RAID-3 or RAID-5 system.

partition — A distinct area of a hard disk that has been prepared to store a particular type of data.

partition table — Information that defines the size and location of each partition on a hard disk.

PATH — An environment variable containing a list of directories on the Linux system that the shell searches each time a command is executed.

Perl — A programming language developed by Larry Wall initially to process strings of text and generate reports.

ping — A command used to test a network connection.

pipe — A connection between two Linux commands (indicated by the | character) that causes the output of one command to be used as the input of a second command.

plain-text configuration file — A file containing human-readable instructions that are used by a program to set its configuration information.

points of failure — Parts of an information system that are subject to failure.

positional variable — A variable used within a shell script that contains data included on the command line when the script was launched.

power supply — The component within a computer system that converts the incoming AC power from a wall socket or UPS device to the correct voltage for use by components in a computer.

preemptive multitasking — A technique used by the Linux kernel to control which program is running from moment to moment.

print filter — A script that contains instructions for formatting documents using the page description language required by a specific printer. The print filter is used by the lpr program to prepare files to be sent to a physical printer.

print job — A file submitted for printing via the lpr command.

print queue — A subdirectory where files are stored to wait for the lpd daemon to retrieve them one by one and send them to the printer. Also called a print spool directory.

print spool directory — *See* print queue.

printcap — The printer definition file used by the Linux printing system. This file specifies how and where files to be printed are stored and processed by lpr and lpd. It stands for *printer capture*.

printer policy — A brief document that describes how print resources can be used and how the management of the printers will be conducted within an organization.

priority — A number indicating the severity of a message submitted for logging. (Log configurations are often based on the priority of incoming messages.) Also, a value assigned to a process running on Linux that determines how much CPU time is granted to the process to complete its tasks.

process — A task running on a Linux operating system, managed by the Linux kernel.

process ID (PID) — A number from 1 to 65,000 that is associated uniquely with a process running on a Linux system.

program — An imprecise term used to refer to any process running on a Linux system.

programming language — *See* computer language.

programming libraries — Collections of programming functionality that are common to many applications. *See also* shared libraries.

Project Athena — The project sponsored by DEC and MIT to create a graphical environment or windowing system for UNIX.

protocol — An organized pattern of signals or words used to communicate efficiently.

ps — Command used to obtain detailed information about processes running on Linux.

pwd — Command that displays the current working directory.

Python — A scripting language developed by Guido van Rossum that is often used for creating graphical programs.

queue — A list of commands or files to be processed.

RAID — A system using multiple inexpensive hard disks arranged in a predefined pattern (an array) to improve performance, increase fault tolerance, or both.

RAID-0 — A RAID level that uses striping to improve disk performance but without adding any fault tolerance.

RAID-1 — A RAID level that uses disk mirroring to significantly improve fault tolerance. Disk read performance is also improved, but disk write performance suffers.

RAID-3 — A RAID level that uses striping with parity information to improve both performance and fault tolerance.

RAID-5 — A RAID level in which data striping with parity is spread across all disks in the RAID array (compared to RAID-3, in which the parity information is stored on a single hard disk).

RAID-Linear — A method of combining multiple physical devices into a single logical device.

RAID subsystem — A hardware-controlled RAID device containing a CPU and other components to control the array of hard disks.

random access memory (RAM) — Electronic storage used by a computer as a working space for all operations while the computer is turned on.

read-only memory (ROM) — Nonvolatile electronic storage within a computer. Used to store information about how the computer starts and how the devices in the computer are configured.

read permission — A file permission that allows a user to read the contents of a file or browse the files in a directory. Represented by a letter `r`.

reboot — Command used to shut down all services and then restart the computer.

redirection — The concept of changing the location where a Linux program receives its input and sends its output in order to increase flexibility and interaction with other Linux programs.

redundancy — Term used to refer to a duplicate system component or piece of data. Many fault tolerant systems rely on the use of redundant components or data; in the event of a failure, the duplicate component or copy of the data would still be available.

redundant arrays of inexpensive disks — *See RAID.*

regular expression — A system of expressing patterns using special characters that can be interpreted by many Linux programs.

regular user account — A user account that, unlike the `root` account, is not used for system administration work. A regular user account has a name similar to a person's name.

release number — A number assigned by the company that prepares a Linux product. It allows the company to track how many times the kernel file has been altered before the final product is shipped.

renice — Linux command used to change the nice level of a Linux process that is already running.

rescue disk — A disk created specifically to boot a Linux system in the event of a system failure. Contains the software tools most likely to be of help in diagnosing and repairing the problem with the failed system.

rescue mode — A mode of operation in Red Hat Linux that is initiated by starting the system using a rescue disk. Rescue mode is used to repair a system failure that blocks normal booting and operation.

resource database — A file that defines how an X resource should appear on screen.

resource groups — The tasks and their accompanying system resources that are defined within a high availability cluster. Each server in the cluster can take over a complete resource group if the server handling that resource group fails.

restore — To copy data from a back-up location (for example, a tape cartridge) onto the file system where that data is normally used, and from which it was unintentionally lost.

rm — Command used to delete a file.

rmdir — Command used to remove (delete) an empty directory.

rmmod — Command used to remove a module from the kernel.

root — Superuser account name in Linux.

root directory — The starting point for all access to Linux resources. It is indicated by a single forward slash: /.

root disk — A disk used in conjunction with an install disk during some Linux installations (depending on the Linux distribution you are using).

rotating log files — The process of moving existing log files to another filename and location for archiving or review. Rotating log files frees hard disk space for new log messages.

rpm — Command used to manage all of the rpm software packages on a Linux system.

run level — A mode of operation that defines which Linux system services are started and which are shut down. The standard Linux run level of 3 includes networking and other common services such as system logging and task scheduling.

safedelete — A type of utility that makes files appear to have been deleted but actually saves a compressed copy of each one in case it is needed later.

SAGE (System Administrators Guild) — A professional organization for system administrators.

script — *See* shell script.

SCSI — A high-performance interface used by many types of devices to connect to a computer.

sector — A unit of data storage on a hard disk. Normally a sector contains 512 bytes.

sed — A command used to process each line in a text file according to a series of commands provided by the user.

selection statement — A statement that lets a computer programmer determine which parts of a program will be executed according to values that are calculated by testing as the program is executed. The *if/then* statement is a selection statement used in shell scripts.

selector — A field in the `syslog.conf` file that determines what events are being logged. The selector is composed of a facility and a priority.

session — A configuration that defines a set of graphical programs to run when a user logs in.

set — Command used to display a list of all environment variables defined in the current environment.

Shadow Password system — Security system used to restrict access to encrypted password text.

shared libraries — Programming libraries that are used by several dynamically linked applications running at the same time. Shared libraries are used to reduce memory consumption by providing a single copy of redundant functionality.

shell — A command-line interpreter, providing a command-line interface.

shell prompt — A set of words or characters indicating that the shell is ready to accept commands at the keyboard.

shell script — An executable file containing lines of text as you would enter them at a command line, including special commands to control the order in which lines in the file are executed.

shell variable — A variable used within a shell script to store information for use by the script.

shutdown — Command used to shut down Linux gracefully.

SIGHUP — A signal sent to a logging daemon to instruct the daemon to reread its configuration files and the log file it writes to.

signal — A message (one of a fixed set determined by the Linux kernel) that can be sent to any process and responded to according to how that program is written.

single point of failure — A system component which, if it alone fails, renders a system unusable. Can be a hardware component or a piece of data.

software — Instructions that control the physical computer components, but which can be changed because they reside on a changeable media such as a hard disk.

software license — A legal definition of who can use a piece of software and how it can be used.

software masters — Original copies of an application supplied by a software vendor or manufacturer, usually one or more CDs, tapes, or disks.

software package — A single file that contains all the files needed to install and use an application or group of related applications. Special data formats are used to store many files in a single software package.

software-based RAID — A RAID array that is controlled or managed by software on the computer system that uses the hard disk array (the Linux system) rather than by a separate CPU or other hardware components.

sort — A command used to sort all of the lines in a text file, writing them out in alphabetical order or according to options provided to the command.

source code — The file that a programmer writes using the keywords and syntax rules of a computer language.

Stallman, Richard — Founder of the Free Software Foundation and the GNU project.

standard error (STDERR) — The channel used by most Linux programs to send information about errors in program execution.

standard input (STDIN) — The communication channel used by most Linux programs to collect input (normally from the keyboard).

standard output (STDOUT) — The communication channel used by most Linux programs to write output (normally to the screen).

statement — A command within a computer program. A statement is often a single keyword, but the term may also refer to a group of keywords that the computer language syntax requires or allows to be used together.

statement block — A list of commands (or statements) that are controlled by a selection or loop statement.

statically linked applications — Linux programs that include library functions in the program itself so that they are not dependent on the libraries loaded on the Linux system.

striping — A technique in which parts of a file are written to more than one disk in order to improve performance. *See RAID-3.*

su — Short for substitute user. Command used to take on the identity of a different user account.

superblock — The master information record for a Linux file system.

superuser — The `root` user account, which has supervisory privileges throughout the Linux system.

surge suppressor — A device that prevents potentially damaging electrical irregularities from reaching a computer system's power supply.

swap partition (also called swap space) — A designated area on a hard disk used as virtual memory by the Linux kernel.

swapon — Command used by Linux initialization scripts to activate the swap partition defined in the `/etc/fstab` file.

symbolic link — A file record that includes a path and filename, but not an inode. It refers to another filename rather than to data in a file.

symmetrical multiprocessing — A technique that allows an operating system to support multiple CPUs on the same computer.

sync mode — An option assigned to log files by which the data written to the log file is immediately written to the hard disk rather than being cached in memory to improve system performance.

syntax — A formalized arrangement of information to allow a Linux command to understand parameters, options, and so forth.

syslog — A programming function used by Linux programs to submit messages for logging. The `syslog` function interacts with the `syslogd` daemon to write messages according to the `syslog.conf` file.

syslog.conf — The configuration file used to control how and where messages are logged by `syslogd` and `klogd`.

syslogd — The background program (or daemon) that manages all of the calls to the `syslog` function, writing log messages according to the `syslog.conf` configuration.

system initialization scripts — Instructions executed each time you boot your Linux system.

tab completion — A feature of the shell that lets you enter part of a file or directory name and have the shell fill in the remainder of the name.

tar — Command used to create a single archive file that contains many other files, often compressed to save space.

tar archive — A file created by the `tar` command.

target hard disk partition — The location on the system's hard disk where Linux will be installed. Also known as *target partition*.

Tcl/Tk — A scripting language developed by John Ousterhout that is used to create graphical programs.

telinit — Command used to switch the system to a different run level.

TENEX/TOPS C shell (TC shell) — An enhancement of the C shell. This is the version of the C shell that is commonly provided on Linux systems.

terminal emulator window — A command-line window (also called a terminal window) within a graphical environment.

test — A command that evaluates the arguments provided after the command name and returns either true (a value of 1) or false (a value of 0).

test — A method of examining data within a shell script and acting according to the result of the examination (or test).

TeX — A document processing system that writers use to create large and complex documents on UNIX or Linux systems.

then — A command that identifies the commands to be executed if the test introduced by the `if` command succeeds (returns a value of true).

thrashing — Excessive movement of processes between RAM and swap space, resulting in reduced system performance and excessive wear on the hard disk.

thread — A piece of a process (or a piece of a daemon, since a daemon is a type of process). The distinction between processes and threads is not important for most system administration work. Instead, the term *process* is used in most cases that don't involve programmers developing software for multiprocessor computers.

timestamp — A record of the date and time when an event occurred.

top — Linux command used to view the most CPU-intensive processes running on Linux at a given moment, along with related process information for those processes.

Torvalds, Linus — Originator of the Linux kernel; formerly a student in Helsinki, Finland.

touch — Command used to create an empty file or to update the access time of an existing file.

tracks — The concentric circles on each platter of a hard disk.

troff — A command used to format and display documents that are created using roff mark-up codes.

tty — A command that displays the name of the terminal device you are currently working in.

tune2fs — A utility used to view or adjust parameters within the superblock of a Linux file system.

umask — Command used to set the file permissions assigned when you create a new file.

umount — Command used to unmount a file system that is accessible as part of the Linux directory structure.

uname — Command used to provide information about the operating system, including the kernel version.

uninterruptible power supply (UPS) — A device capable of providing power to a computer via batteries when the incoming AC power (wall socket power) fails. It also informs the Linux system of the status of the power.

UNIX — An operating system created at AT&T Bell Labs (now part of Lucent Technologies) about 30 years ago by Ken Thompson and Dennis Ritchie. UNIX is still widely used, and it provided the technical basis for Linux.

user permissions — A set of three file permissions (r, w, and x) that apply to the owner of a file or directory.

User Private Group — Security system that creates a new group containing one user when that user is first created.

useradd — Command used to create (add) a new user account in Linux.

usermod — Command used to modify or update an existing user account.

utility programs — Software that provides assistance in managing the hardware and operating system features (as opposed to doing other types of work such as word processing).

variable — A memory location used by a program to store a value, such as a number or a word. Each variable is assigned a name so that the program can access the value by referring to the name.

virtual memory — Memory available to the Linux kernel for running programs but which is actually located on a hard disk. Data that the Linux kernel stores in virtual memory is placed in the swap file system, or swap space.

vmstat — Linux command used to display detailed information about virtual memory usage.

Webmaster — The person who manages the content and functioning of a Web server program running on Linux.

wheel — Special system administrative group, not used officially in Linux.

while — A command that creates a loop based on a test. The loop executes a block of commands as long as the test returns true.

white space — Tabs or spaces included in a program or script that make the script easier for a person to read.

window manager — A special-purpose graphical application (X client) that controls the position and manipulation of the windows within a graphical user interface.

Windows NT — A business-oriented operating system product developed by Microsoft. Windows NT is not based on DOS as an underlying operating system.

write caching — A technique in which information to be written to a file system (particularly a RAID file system) is stored in system memory temporarily in order to improve performance of reading and writing information to the file system.

write permission — A file permission that allows a user to add or change information in a file or create files within a directory. Represented by a letter w.

WYSIWYG — A characteristic of programs that show documents on the computer screen much as they will look when printed on paper or in a Web browser (pronounced "whiz-ee-wig").

X client — A graphical application.

XFree86 Project — A free software project that creates software to provide X Window System functionality to Linux.

X resource — The separate screen elements of a graphical application, such as scroll bars, text fonts, mouse pointers, and title bars for windows or dialog boxes.

X server — The program that communicates with the video card to create images on the screen.

X Window System — A graphical software environment used by almost all UNIX and Linux operating systems.

xconfigurator — A utility in Red Hat Linux for configuring the X Window System.

xdvi — Program used to display a LaTeX-coded document as it will appear on paper.

xf86config — A standard text-based utility for configuring the X Window System.

xf86Setup — A graphical utility for configuring the X Window System.

xfontsel — Program that lets the user choose each aspect of a font definition (such as the font family and typeface) and then displays the corresponding font for review.

xload — A graphical program that displays the CPU load over time on any Linux system's graphical interface.

xrdb — A command that loads an initial X database resource file or merges additional resource configuration details.

xterm — A program within a graphical environment that provides a command-line window.

YAST — A graphical configuration utility developed by the makers of SuSE Linux.

zcat — Command used to print the contents of a compressed file to the screen.